REMARKS/ARGUMENTS

Claims 1-22 are pending. Claims 1, 2, 8-11, 15, and 18 have been amended. New claims 21 and 22 has been added. No new matter has been introduced. Applicants believe the claims comply with 35 U.S.C. § 112.

Claims 1-14

Claims 1-14 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over Bridge (US 6,530,035).

The present application relates to storage systems, and in particular to storage system management in which failure boundaries are taken into consideration when assigning storage volumes. Specifically, failure boundaries are determined to ensure that replication storage volumes cross failure boundaries so that the impact of a failure event within a given failure boundary may be minimized.

Applicants respectfully submit that independent claim 1 as amended is patentable over Bridge because, for instance, Bridge does not teach or suggest determining a boundary of a failure of the primary storage volumes and the replication storage volumes, the boundary being determined using error correction group and controller group information of the primary storage volumes and replication storage volumes to divide the storage volumes into failure groups of logical volumes.

Applicants respectfully submit that independent claim 10 as amended is patentable over Bridge because, for instance, Bridge does not teach or suggest a memory for storing information regarding at least one boundary of a failure of the primary storage volumes and the replication storage volumes, the at least one boundary being determined using error correction group and controller group information for the set of primary storage volumes and the set of replication storage volumes to divide the storage volumes into failure groups of logical volumes.

The error correction groups include one or more of the logical volumes used within the storage system. As further stated in ¶[0029] of the present application, "[t]he failure boundary is used to designate that if one of the disk drives 121 in the ECC group 301

fails, all of the other logical volumes which belong to that group 301 will be impacted." Fig. 3 further illustrates this point in showing that the ECC groups are made up of volumes:

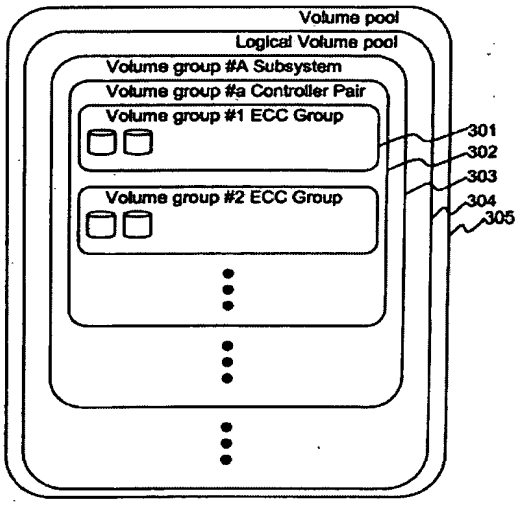

In contrast, the failure groups within Bridge are divided amongst disk drives, not logical volumes. As described in column 13, lines 32-56, each disk drive can be associated with a failure group where they share some common failure criteria. Bridge specifically states in regard to Fig. 6: "Disk drives 624, 626, 628, and 630 share a common projected failure condition which results in these disk drives being assigned to the same failure group 622. Thus, disk drives are in the same failure group if there is a failure mode that could affect each of the disk drives and redundant data is maintained to protect against that failure." Furthermore, Fig. 6 illustrates three embodiments of the present invention in which disk drives are assigned as failure groups, instead of logical volumes:

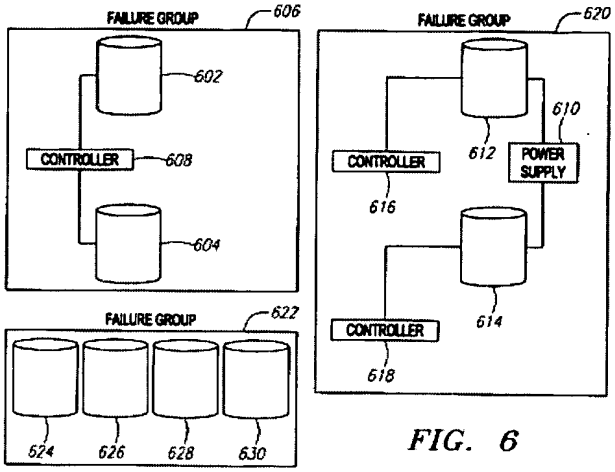

FIG. 6

In addition, the parity protection method utilized within Bridge utilizes the mirroring of data extents or contiguous storage sections on disks rather than separating failure boundaries based on error correction groups. Nothing in Bridge teaches or suggests dividing storage volumes into *failure groups of logical volumes by a boundary of failure determined using error correction group and controller group information.*

For at least the foregoing reasons, Applicants respectfully assert that independent claims 1 and 10, and dependent claims 2-9 and 11-14, are novel and patentable over Bridge.

Claims 15, 16, 18, and 19

Claims 15, 16, 18, and 19 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over Bridge in view of Wahl et al. (US 6,618,818). The Examiner cites Wahl et al. for allegedly disclosing determining the boundary of a potential failure based on logical addresses.

The Examiner asserts that it would have been obvious to combine Wahl et al. with Bridge to render the claims obvious. According to the Examiner, the motivation to divide the failure groups into logical groups is that Wahl et al. discloses grouping storage devices into logical groups. That is not an adequate basis for combining Wahl et al. with Bridge.

Significantly, Bridge addresses a specific problem, i.e., the "multiple-disk failure problem." "The solution to the multiple-disk failure problem is to limit the number of disk drives that are used to store redundant data. Thus if disk drive X fails, a second failure will not lose data unless the failure occurs to one of N other disk drives where N is a relatively small number. To accomplish this in an embodiment of the present invention, each named disk drive can be associated with a 'failure group.'" Column 13, lines 28-34. "All named drives in a failure group share some common disk drive failure criteria, which is any failure mode or condition which is projected to cause the related disk drives to fail at the same time period." Column 13, lines 35-38. Therefore, Bridge discloses failure groups divided amongst disk drives wherein the disk drives in a failure group share some common disk drive failure criteria. There is no reason to change the division amongst disk drives into the division amongst logical volumes because the latter is not designed to overcome the *multiple-disk failure problem* using a technique that relies on the *common disk drive failure criteria*.

Moreover, the present invention divides the storage volumes into failure groups of logical volumes using error correction group and controller group information to set the boundary of failure. This is not taught in Bridge or in Wahl et al. Bridge discloses that the failure groups are divided amongst disk drives according to common disk drive failure criteria. Wahl et al. merely discloses that storage devices can be grouped into logical groups. Both references are devoid of any teaching of using error correction group and controller group information to set the boundary of failure. Even if combined, Bridge and Wahl et al. would not render the claims obvious.

For at least the foregoing reasons, claims 15, 16, 18,and 19 are patentable.

Claims 17 and 20

Claims 17 and 20 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over Bridge in view of Tam (US 6,411,969). The Examiner cites Tam for allegedly disclosing that the replication process utilizes a daily or hybrid backup implementation.

Even assuming that Tam supplies the missing teaching, Claims 17 and 20 are still patentable because Tam does not cure the deficiencies of Bridge, in that Tam also fails to teach or suggest determining a boundary of a failure of the primary storage volumes and the replication storage volumes, using error correction group and controller group information of the primary storage volumes and replication storage volumes to divide the storage volumes into failure groups of logical volumes. For at least the foregoing reasons, claims 17 and 20 are patentable.

New Claims 21 and 22

New claims 21 and 22 depend from claims 1 and 10, and further recite that the primary storage volumes and replication storage volumes are horizontally or are vertically addressed. These features are also absent from Bridge. Therefore, dependent claims 21 and 22 are patentable over the cited art.

CONCLUSION

In view of the foregoing, Applicants believe all claims now pending in this Application are in condition for allowance and an action to that end is respectfully requested.

If the Examiner believes a telephone conference would expedite prosecution of this application, please telephone the undersigned at 650-326-2400.

Respectfully submitted,

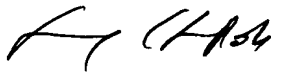

Chun-Pok Leung
Reg. No. 41,405

TOWNSEND and TOWNSEND and CREW LLP
Two Embarcadero Center, Eighth Floor
San Francisco, California 94111-3834
Tel: 650-326-2400
Fax: 415-576-0300
RL:rl
60833520 v1